

A configurable accelerator for manycores: the Explicitly Many-Processor Approach

János Végh

Abstract—A new approach to designing processor accelerators is presented. A new computing model and a special kind of accelerator with dynamic (end-user programmable) architecture is suggested. The new model considers a processor, in which a newly introduced supervisor layer coordinates the job of the cores. The cores have the ability (based on the parallelization information provided by the compiler, and using the help of the supervisor) to outsource part of the job they received to some neighbouring core. The introduced changes essentially and advantageously modify the architecture and operation of the computing systems. The computing throughput drastically increases, the efficiency of the technological implementation (computing performance per logic gates) increases, the non-payload activity for using operating system services decreases, the real-time behavior changes advantageously, and connecting accelerators to the processor greatly simplifies. Here only some details of the architecture and operation of the processor are discussed, the rest is described elsewhere.

Index Terms—computer architecture, processor accelerator, manycore processor, many-processor approach



1 INTRODUCTION

ABOUT a decade ago, “growth in single-processor performance has stalled – or at best is being increased only marginally over time” [1]. Since the computing industry heavily influences virtually all industrial segments, the stalling crashed the long-term forecasts in the industry, economy, etc. The need for growing performance triggered research in many directions, from “rebooting computing” [2] through using a “cross-layer approach” [3] to using accelerators combining different technologies, like reconfigurable General-Purpose Graphics Processing Unit (GPGPU) [4]. Especially the popular and powerful many-core processors provide huge computing capacity and sever problems, like “multi-core and manycore vendors and runtime systems cannot possibly support the virtually unlimited number of processor configurations” [5]. It is not yet checked, however, what hidden reserves for acceleration can be disclosed in the operation of the conventional processors itself.

In addition to the hardware (HW) issues, also the software (SW) makes its contribution: “parallel programs ... are notoriously difficult to write, test, analyze, debug, and verify, much more so than the sequential versions” [6], and doubts like “Chip multi-processors have emerged as one of the most effective uses of the huge number of transistors available today and in the future, but questions remain as to the best way to leverage Chip Multi Processor (CMP)s to accelerate single threaded applications” [7]. This is why “cross-layer approach spanning from hardware to user-facing software is necessary to successfully address this problem” [8].

In this paper a “better way” of accelerating single threaded applications is searched, first scrutinizing the hidden reserves in the operation of the stored program processor based computing systems. During this, it is shown that

the final reason of the present stalling (in addition to the already known reason: the finite speed of the light) is the 70-years old single-processor approach, which dominates both computer construction and programming. The introduced Explicitly Many-Processor Approach (EMPA) viewpoint solves some old problems of computing. It increases the computing throughput, depending on the context up to several dozens times higher, allows to build more deterministic real-time systems, etc. At the same time, it allows to simplify the internal architecture, to use less transistors for the chips.

2 THE SINGLE-PROCESSOR APPROACH

According to the state of the art of his time, Neumann considered only one processor and formulated the principles of operating a processor considering only the execution of a single machine instruction. Considering his paradigms carefully, neither of them contradicts to the present many-processor approach (see section 3).

2.1 Lack of time dependence

The roots of the event-oriented processing were implicitly present in Neumann’s approach: the next instruction can only be executed when the processor is not any more busy with executing the current instruction. The time, however, in the paradigms is considered as an implementation detail. Although later the lack of considering the time explicitly was identified as one of the fundamental issues in computing [9], handling time (as well as synchronizing) became the task of operating systems, which they can solve in a quite resource-wasting way [10].

The availability of the processing unit, the instruction and the data are critical factors, and all they have a timely behavior. The most successful approaches to improving performance of processors modified the conditions of processor

• J. Végh is with Faculty of Mechanical Engineering and Informatics, University of Miskolc, Hungary
E-mail: J.Vegh@uni-miskolc.hu

Manuscript received July 10, 2016; revised August 26, 2016.

availability (methods for reducing the instruction and data memory access time are not touched here). The pipelining separated the signal into "ready to accept data" and "result is available" signals, the hyper-threading connected the ready-to-run thread to the processing unit. Critics like *"HT generally improves processor resource utilization efficiency, but does not necessarily translate into overall application performance gain"* [11], call to scrutinize the processor availability condition and cross-layer functionality, if one wants to improve single-processor performance using several processing units.

2.2 Atomic processing unit

Hyper-threading separates the hardware accelerators (like registers and core-level cache) from the raw processing power, which is a clear recognition of the fact that from programming point of view, the single machine instruction (a HW unit) is too small (cannot make reasonable functionality), while a complete process (a SW unit) is too big (wastes time with waiting), so some intermediate unit: HW-supported Quasi-Thread (QT) should be introduced. This suggests to check, *what the optimum size of the right unit is, and how it should be supported from HW and SW*. Also it shows that *the effective problem solution can only be reached through HW/SW codesign, i.e. through crossing HW/SW layers* [8].

2.3 Multiple processing units

From again another side, from the simple out-of-order processing through multiple arithmetic units to speculative evaluation, several kinds of *hidden* processing units have been introduced and they successfully accelerated the operation of the processor. There is no question, that more processing units are needed to make parallel operations (and in this way apparently the processing) quicker. However, those solutions forget about one of the most important principles of Neumann: *computer should be simple*.

In order to achieve higher performance, *"computers have thus far achieved this goal at the expense of tremendous hardware complexity – a complexity that has grown so large as to challenge the industry's ability to deliver ever-higher performance" ... and ... "the ever-increasing complexity of superscalar processors would have a negative impact upon their clock rate, eventually leading to a leveling off of the rate of increase in microprocessor performance"* [12]. This suggests to check *what the optimum way of using multiple processing units is, without wasting computing and electric power?*

2.4 Multitasking issues

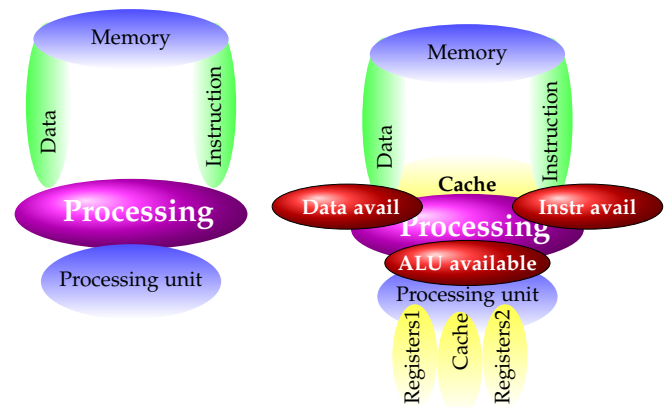
Different reasons directed the designers to share the computing capacity and other resources between different tasks, starting with the age of single-processor systems. The methods, however, have been prolonged to the era, when several processing units could be used. Even today, the external peripherals *interrupt* the control flow (although some other core could do interrupt servicing), in the many-processor systems all Processing Unit (PU)s are *central*, the operating systems providing services and scheduling the operation of the running tasks take the processor time from the payload jobs, changing context between user and kernel modes causes a considerable non-payload activity, the hardware

scheduling makes the software operation non predictable, etc.

A related special issue for accelerators is that an accelerator is always outside the processor, and it is efficient because it works differently from the conventional, programmable processors. Several problems must be solved in order to connect the stored program (von Neumann) processors with the rest of the world. The processor only offers input/output (I/O) bus for connecting an accelerator to the processor. However, in the today's environment an Operating System (OS) must provide protection (virtualization) for the I/O operations. It is only possible in protected mode, and the context change from and back to user mode is extremely expensive: it is in the range of dozens of thousands clock periods for the modern HW architectures and OSs [13]. This fact *restricts the utilization of general-purpose accelerators to accelerate only activities long enough to be not disproportional with that offset time*.

3 THE MANY-PROCESSOR APPROACH

The principle of operation of stored program processors (see Fig 1(a)) and its engineering implementation (see Fig 1(b)) became quite different during the time, due to the efforts to enhance processing speed of the computer. Some of the enhancements are none-essential: although without cache memories the operation of the processor would be painfully slow, the processor would work. Also, the highly successful accelerators, the internal registers, are not strictly needed for the operation. The processors could work with a strongly limited number of registers [14], or as the example of the very first EDVACs prove [15], even without registers. Also recall, that it is advantageous to separate registers (as "glue" material, together with internal state and cache) from the processing unit, see hyper-threading or shadow register set at some interrupt servicing.



(a) Principle of the computation in stored program computers (b) Engineering implementation of stored program computers

Fig. 1. Comparing the theoretical and engineering operation of computers

The many-core processors provide an ideal field to implement an EMPA environment. The idea here is to *provide the cores with the ability to outsource part of the job they received*. I.e. they remain responsible for doing the job, but not necessarily "personally". The approach mathematically covered

by the theory of “communicating serial processes” [16] and some additional aspects are discussed in [17].

It requires of course a lot of changes: some external logic must concert the work of cores (data transfer, dependency and synchronization), the information about the possible outsourcing must be prepared at compile time rather than at runtime, the code must be cut to optimally sized, partly independent QTs, the processor must be notified about the pre-calculated parallelization possibilities, etc. The QTs can be embedded into each other, so the core receiving some outsourced QT can also attempt to find helper cores¹ to make further outsourcing.

3.1 Availability of processing unit

In the Neumann paradigms, the task of the control unit is only to provide a “proper sequencing”. In the single-processor approach, the control unit allows sequential stepping, jumping, calling and returning, and even interrupt servicing.

In many-processor approach, it is also possible to report the availability of several processing units (this time explicitly, unlike hidden processing units in the single-processor approach). For compatibility, a new control layer is introduced. The supervisor (SV) is a second control layer in the many-core processor above the cores, which (among other functionalities) partly takes care of providing this availability signal. Since it knows about all cores in the processor, it can provide ‘ready’ signal as long as at least one core is ready to process instructions. In this way the processor is able to receive new instructions as long as at least one of the cores is ready to work.

3.2 Atomic units

Unfortunately, only absolutely independent machine instructions can be distributed between PUs in this way: the “glue” carries valuable state information between individual machine instructions. For allowing this internal state transfer, the code must be cut into QTs in a reasonable way. However, to select the proper *granularity* is hard.

As mentioned, the machine instructions are (in general) too small to form a QT. The ideal way would be to use the PUs as stateless automatons, which would result in too coarse granularity, like threads. To allow for a more effective QT size, a *quasi-stateless automaton model* is used. Upon beginning to execute a QT, a full-fledged “glue” set, comprising register file, flags (clone) and cache (shared), is provided for the newly hired PU, and a limited amount of “glue” can be returned in a (by the supervisor) synchronized way when a QT is finished. Since QTs are running on the cores in a separated way, for providing atomic units, using QTs is ideal: when the core returns ‘ready’ signal, the action protected by the QT is finished: both ‘owner’ and ‘others’ must wait it.

1. Since in any given moment there is a one-to-one correspondance between an allocated core and the QT running on it, these two terms are used interchangeably: they emphasize the HW and the SW side of the same thing, similarly as the processor and the process do in single-processor approach.

3.3 Mapping QTs to cores

As the QTs can create other (child) QTs, they form a kind of processing *graph*. A virtually infinite number of graph nodes must be mapped to a finite number of physical cores, so sometimes the new QTs must wait for computing resource. The supervisor’s core allocation algorithm prefers reaching the leaves of the graph rather than opening new forks, and as an emergency mechanism for the case when SV runs out of available cores, the cores can suspend processing their own QTs, borrowing their own resources to their child-QTs while they are executed.

Another crucial question is how the at compile time labelled QTs can refer to the QTs distributed at run time to cores, having “random” availability. The QTs wanting to communicate with other QTs refer to the compile-time address of the QT and the SV translates this address to a runtime physical core number. Because of this, QTs have data structures similar to the processes running under OSs.

3.4 Synchronizing cores

The supervisor maintains configurable² parent-child relationships between the cores, so the cores have information whether a QT is running, either in their own scope, or in scope of their parent. Without explicit synchronization, the cores work in parallel on different QTs of the originally single-threaded code. This kind of synchronization is done at processor level, in one clock cycle; rather than the conventional methods, which need awfully long times [18]. Even, since waiting is handled by the SV based on signals handled by the SV, no time is used when there is no need to wait.

3.5 Linking cores

As contented by Amdahl [19], to increase further the performance, some cooperation between cores is necessary. In EMPA it is implemented in such a way that the cores pass the control signals and data to the SV, which – like a telephone switching center the calls – forwards them to the right core. The SV plays a key role here: it handles all resources of the processor. It connects and synchronizes the cores and delivers the limited amount of information, and even, it provides partners when one of the cores signals that it could advantageously use more processing capacity. Also, it concerts collective processing and takes over some functionalities from the OS, which functionalities are traditionally missing from processors.

When creating a new child QT, the new core must be able to continue the processing from the point where its parent is, i.e. the “glue” of the parent must be *cloned* (using dedicated wiring) to the child. Also, upon termination, the child might send back the content of the link register (clone back, using similar method). The cores are working in parallel, so a special synchronization mechanism must be used. The data intended to be sent to the other party, are latched by the sender (see also Figs. 2 and 3). Later some triggering signal transfers the data to its destination, where they are latched

2. This feature is between programmable, since it takes the configuration information from the object file and reconfigurable, because it changes the type of connection between the components inside the processor

again, and the cores need explicitly use the latch register to access its content. In this way all partners send data to and receive data from their partner when they need it, providing a proper synchronization. The cores see those transfers as using pseudo-registers (see section 4.6).

3.6 Subroutines, interrupts, traps

A main task related to subroutines is to remember the return address. In the EMPA way, starting a QT means providing also a new core for the QT, and the processing in the parent code continues immediately, at the address next to the QT, i.e. at the return address. That means, that the "return address" shall only be stored for the time of creating a QT (one clock period) and it is automatically remembered by the SV. So, a QT behaves very similarly to a subroutine call, but the QT itself is embedded in the "calling" code flow. Implementing a special metainstruction for subroutine call just allows to place the body of the subroutine outside the main code flow.

In the single-processor approach, when another process is to be run, the processor must be stolen from the running process, which needs a lot of saving and restoring, as well as context changing in advanced systems; i.e. rather much non-payload activity. In EMPA approach, a core can be reserved for interrupt servicing. It can be prepared (even in kernel mode) and waiting for the interrupt. When the interrupt arrives to the core waiting in power economy mode, it immediately starts its servicing, without any duty to save and restore, saving processing time and memory cycles. Note that *in EMPA approach no context change is needed, resulting in several hundreds of performance gain* relative to the conventional handling.

Similarly, cores can be prepared to provide kernel services. Some system services, for example semaphore handling, do not really need all the facilities of the OS, they can be implemented in some alternative way. As our former measurements on soft system [20] proved, such alternative implementation resulted in performance gain about 30, although in that case no context changing was needed. Similar gain can be expected when implementing OS services with EMPA. The gain factor will surely be increased because of the eliminated context change, but the concrete gain will depend on the functionality of the service. In addition to making context changing (in both directions) obsolete, the kernel and user codes can run even partly parallel.

3.7 Mass processing

The physical vicinity also allows to implement certain kinds of cooperation between the cores. Here the source of performance gain can be to eliminate control-type instructions (like loop counter advancing, checking, jumping) as well as to eliminate unnecessary stages of instruction execution, like read and writeback when only the final result is interesting.

A typical example is summing up elements of a vector (see Listing 1): the read and write back of the partial sum is needed only when considering the machine instructions as atomic unit. In mass operating mode, the parent can sum up summands provided by its children, in frame of a machine instruction.

3.8 Linking special accelerator

As it follows from the description above, for the SV a core is represented as a source and destination of signals and data. i.e. an extremely simple interface is provided for linking QTs. Since the QT can be as large as a SW thread, and SV only "sees" the signals and data, but no HW at all, EMPA provides an extremely simple interface for linking any kind of external accelerator (without any non-payload activity) or even HW-implemented SW processes like [21].

4 ARCHITECTURE OF AN EMPA PROCESSOR

Many of the experts of the field expect some kind of solution from the reconfigurable (RC) technology (see for example [22] and its cited references). However, since the computing density [23] cannot be increased using this technology, the RC elements are mostly used to control the conventional components of the processor (for example [4], [24]) or to accelerate some OS services [25], or one can find even complete real time (RT) OSs (for example [26]) implemented in HW. At the other end of the scale one finds also operating systems with SW threads implemented in HW [21], too. When thinking about the EMPA architecture, one should have in mind also RC technology combined with the conventional processor technology. An EMPA processor is akin an end-user configured dynamic (processor) architecture using dedicated wiring and special logic unit (LU)s (cores).

4.1 Using multiple processing units, explicitly

Of course, the introduced new operating principle requires modifications also to the internals of the processor. Since our goal (in addition to boost single-thread performance) is to preserve as much compatibility as possible, with the conventional computing, some components are greatly similar to the conventional ones, and also the newly introduced components have functionalities, which can be formulated using terms similar to the conventional computing.

4.1.1 The processor

For the outside world, the processor is nearly unchanged. It receives a stream of instruction codes, which is enriched with metainstructions, describing the suggestions of the compiler about increasing performance. Internally, however, it works differently.

4.1.2 The cores

The cores in an EMPA processor are mostly similar to the present single-core processor, with some extra functionality. The cores provide and receive extra signals (see also see also Fig. 2) for/from the SV

- **Availability** a core is available when it is not executing a code chunk, not preallocated for a future task, and not disabled for some reason (like overheating)
- **Enabling** when a core is rented, it gets enabled and disabled again when the code chunk terminates
- **Waiting** a core can be blocked when waiting for the termination of some QT running on another core
- **Metainstruction** signals that during pre-fetch a metainstruction found

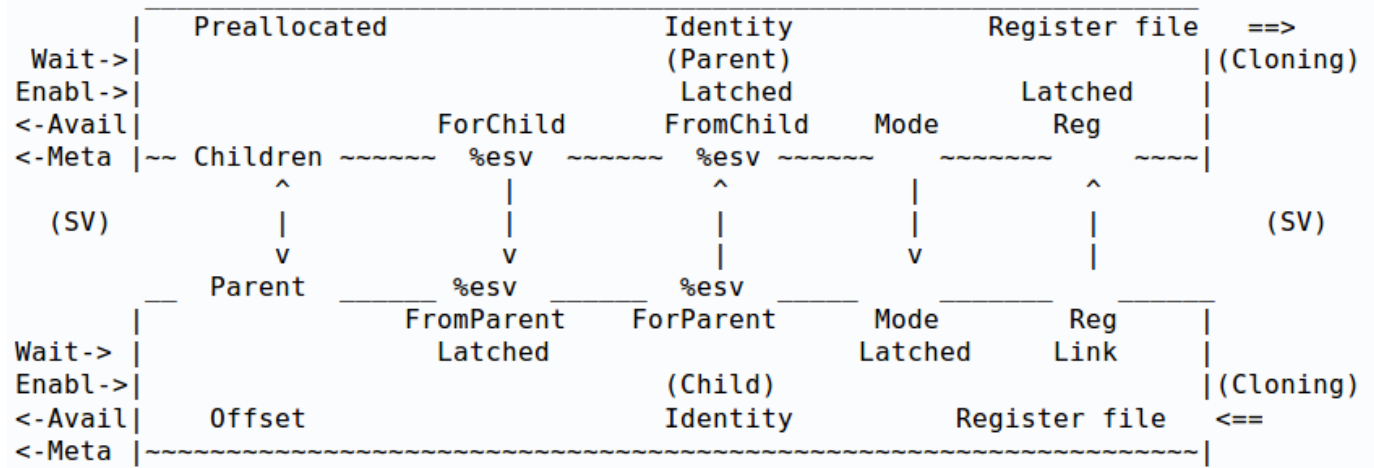


Fig. 2. The communication signals and data between parent and child cores in EMPA

Also store the cores extra information (see also Fig. 2 and section 4.2):

- **Identity** The cores are identified by a (hard) "one-hot" bitmask
- **Parent** The (configurable) identifying bit mask of the parent of the core
- **Children** The (configurable) ORed value of the bitmasks of cores with QT created by the QT running on this core
- **Preallocated** The (configurable) ORed value of the bitmasks of cores preallocated for this core
- **Offset** The (configurable) memory address of the QT the core runs

4.1.3 The supervisor

This new unit is a second, end-user configurable control layer, which, among others, handles the computing resources and reports "ALU avail" signal. It can be considered as a kind of second-level control unit, implemented as a control layer above and between the PUs (cores). It transfers signals to and from cores; as such, all data and control transfers can be implemented between SV and PU, so a 'star' topology shall be implemented, and there is no need to wire the cores with all other cores.

The SV is responsible for providing a new processing unit (to rent a core) when needed, to handle properly the termination signals from the rented cores (including handling the mentioned bitmasks), to provide synchronization signals for a core running a QT, provide (limited) synchronized data connection between the cores, etc. It comprises several simple, easy to implement and quick to execute functionalities.

The SV is responsible for all resources, so it can only be used in a sequential way, one operation at a time. It could be a bottleneck for the performable operation of the processor, so its proper handling requires special attention when designing cooperation between the PUs, working in parallel.

The conclusion that "the ever-increasing complexity of superscalar processors would have a negative impact upon their clock rate, eventually leading to a leveling off of the rate of increase

in microprocessor performance" [12] also means, that when introducing a new control layer (SV) on top of cores in the processor, its simple combinational logic can be operated at a frequency, allowing high-speed coordination of the cores' work; much higher than the clock frequency needed for the cores for making sophisticated computations.

4.1.4 The memory

The more PUs obviously need broader memory access bandwidth, but the burden of memory bandwidth is not as bad as one might think for the first look. In the Single Processor Approach (SPA) systems one processor is linked through one bus to one memory decoder. However, in the hyper-threaded architectures several outstanding memory access requests can coexist.

One must remember, the producer/consumer model: the "memory wall" is still active. Even with EMPA, the many-core processors cannot receive more memory contents, than the memory subsystem can produce. EMPA can make, however, good use of multiple memory access devices. This ability might need to change the memory access architecture: the many-core processors might need more than one independent memory buses, the buses can be (time or space) multiplexed, and the memories might need multiple decoders to the same memory address space. To broaden the memory-access bandwidth, independent multi-port memories are needed, like [27].

In many-processor systems a lot of efforts are needed to provide coherent operation of the Central Processing Unit (CPU)s. In EMPA, the PUs have *coordinated operation*, so the accidental simultaneous access can be eliminated by the compiler/SV. For examples see the sample programs below; the interrupt or OS service operation or direct memory access: the logic of the (cooperative) operation excludes the simultaneous access, so using the (relatively slow and energy wasting) shared memories (like in [28]) are not necessary here.

4.2 The parent-child relationship

The cores are uniform and independent, but using the mentioned bitmasks, they can be in parent-child relationship,

to arbitrary depth. A core can have only one parent, but an arbitrary number of children. This relationship allows for several generations, unlike the master-slave relationship, used in some other architectures like the “Desktop Super-computer” [29] and allows for dynamic behaviour similar to that of the “Invasive Computing” [30].

Fig. 2 attempts to summarize the signal and data traffic of the cores. At the top of figure the core is in role *parent* and at the bottom in role *child*. These roles are of course context-dependent, in another context a child can be parent of another core, or a core cannot have a child at all. The shown storages and signals are typical for that role. Akin in Field Programmable Gate Array (FPGA)s, some well-defined, mostly fixed functionality blocks are placed in close vicinity to each other, and the end-user has the possibility to connect them, changing some configuration parameters, which selects one of the predefined functionalities.

4.3 The dynamic architecture

The individual cores take the responsibility for executing a dedicated QT. The QTs can be nested, i.e. a core may face the task to delegate part of its job to another PU. It can be solved (with the active help of SV) using the parent-child architecture. If there is at least one available core, the SV rents it from the pool of cores for the requesting core, and administers it as a child of the requesting core, in both cores in the configurable bitmasks. This means, that the “processing graph” will be mapped to the available cores.

The child core gets enabled and begins its independent (and parallel!) life. A child core might find a termination metainstruction (in contrast with the conventional processor, where only ‘halt’ is possible), which leads to notifying the SV (see also Figs. 2 and 3). The SV administers the termination of the parent-child relationship, and puts back the (former child) core into the pool. From that moment that core might be rented for another task.

The parent is responsible for performing the complete task of the QT it received, even if it delegated part of the job (in form of child QTs) to its child cores. This means, that a parent QT must wait the termination of all of its child QTs, in order to be sure the work completed. To do so, the SV will block the termination of a parent QT until its children mask gets cleared.

4.4 Data passing

A crucial question is passing data between cooperating cores during processing. Splitting the code into QTs in a reasonable way, allows to make a bargain between losing performance because of transferring data and gaining performance because of using more PUs. Anyhow, *some* data passing is inevitable.

When a piece of the code delegated originally to the parent core is delegated to the child core, the child core must inherit also the internal status of the parent, including the contents of the registers. It needs dedicated wiring between the cores and the SV, and (depending on the physical location of the cores) can take somewhat longer time than the other SV operations. In this case the synchronization is not a problem: the child core commences its life *after* it received the needed data.




Fig. 3. How EMPA processor operates at supervisor and core levels

Upon termination of QTs, however, synchronization of the eventual returned data might be an issue: the children cannot know *when* to copy back data into registers of the parent. To solve this problem, the SV latches the data returned by the child, and when the parent is about to terminate or explicitly waits for the termination of the child, transfers it to the parent core. This type of information transfer requires dedicated wiring between the core and the SV, and should be implemented as a two-stage transfer. In this case the SV acts as a switching center, so making dedicated wiring between a core and the rest of cores can be avoided.

4.5 Two-level operation

Fig. 3 is somewhat similar to a state diagram, and might help two understand the operation of EMPA. Actually, the actions take place at two different levels.

At the beginning, the SV “creates” the cores, i.e. it initializes its internal data structures and places them in a pool of sharable PUs. When in the pool, the operation of the cores is of course not enabled. One of the cores gets allocated and will be enabled.

In that state, the core will work as a traditional processor, with the exception that the SV can disable its operation, and during the pre-fetch stage it decides whether the next instruction is a normal executable instruction or a meta-instruction. In the former case, the core executes the executable instruction as the conventional processors do. In the latter case, using its ‘Meta’ signal, the core notifies SV. In response, the SV takes over the execution of the

metainstruction: advances the Program Counter (PC) of the core to the next instruction at the core level, and 'executes' the meta-instruction at the supervisor level.

If the metainstruction is to create a QT, it means that new core(s) must be rented from the pool (the HW is provided), and equipped with proper internals (the SW is provided). Then the new core begins its independent life. When it reaches a metainstruction, it notifies the SV.

Waiting can occur only as an effect of a metainstruction. When there are no more cores in the pool at the moment, or a parent is about to terminate without its children being terminated or an explicit waiting requested, the SV simply disables the core, until the condition fulfilled.

Parameter passing also happens under SV control. The latched registers (to/from child/parent) are filled in the partners correspondingly, when creating or terminating a QT (executing `QxCreate` or `QTerm`) triggers the action. As shown, SV is fully responsible for synchronized data transfer between latched registers, but its operation is based on how the programmer configures SV through issuing the corresponding metainstructions.

4.6 Pseudo registers

Another crucial question is how such unusual data passing can be used ("programmed") in a way, close to the conventional one. The conventional cores use registers for quick manipulation of data, so a useful idea is to use pseudo-registers for transferring data between a child and its parent. In this way both the parent and the child see a "register", which is a well-known term for the cores. However, it is not an item in the register file: it has a register address, but it also has a context-dependent functionality. In this way the parent and the child can share some data in a way, which appears as handling registers.

This pseudo-register might have a bit longer access time (depending on the physical position of the cores and the additional electronic functionality, hidden under the facades of the register and depending on the operating mode), but surely shorter than reaching any memory or using any kind of internal network. Different operating modes for collective work can be defined, and through those (pseudo)registers the quickest possible data transfer can be reached.

In normal mode a pseudoregister behaves (nearly) as a traditional register, except that it is mapped to a latch register of the core. In mass processing mode a pseudo register behaves in a quite special way. For example, what the parent writes in its own register, children can read from their own register, and what children write in their own register, the parent can read from their own register. In order to avoid synchronization issues, the SV latches the sent data when the sender is ready to send it, and allows the receiver to read the data from the latch when the receiver is ready to accept it.

As shown in Figs. 2 and 3, (pseudo)register transfer might happen in several situations, and in all cases under the control of the SV. The first data transfer occurs when renting a new core. First of all, the SV handles the corresponding bitmasks 'Parent' and 'Children', and clones the complete internal state (including the register file and the PC) of the parent to the new child. When ready, the SV enables the child, which begins its independent (and parallel!) life with executing the delegated code chunk.

After this, the parent skips that (logically already executed) code fragment and continues execution at that new address. When the child is about to terminate, it notifies its SV with the 'Meta' signal. In response, the SV latches the content of the link register of the child core for the parent, changes the contents of register 'Children' and 'Parent' correspondingly. The latched register content is available for the destination parent core through issuing an explicit or implicit 'Wait' (recall that termination implies a 'Wait'), and will be written from the latch into the corresponding register only when the parent requests so. In this way no synchronization issues might happen.

When cores are allocated for mass processing, registers 'ForChild' and 'FromChild' will be initialized in the parent core. The parent can write 'ForChild' through writing its own pseudoregister and read 'FromChild' through reading its own pseudoregister. (Remember, the PC of the parent might stall at the address where mass processing begins; and also that SV can read/write any of the registers, independently of the operation of the cores.)

In mass processing mode, the 'Mode' code and 'ForChild' are latched by the child when preallocating cores. If a child writes to its own pseudoregister, the value will be latched in 'FromChild' in the parent, and the next (repeated) QT creation might consider that value. Depending on the 'Mode', the parent might watch and read through reading its own pseudoregister the latched value written by the child to its own pseudoregister. Since for reading and writing pseudoregister as 'parent' or as 'child' there are two different directions (two different latched registers), some special rules determine its context-dependent utilization, see [31]. To forward data (i.e. to transmit, data received on its input, to its output) the core needs to use an explicit copying from the input pseudoregister to the output pseudoregister instruction. For the programming implementation see [31].

5 UTILIZING THE ARCHITECTURE

After having an architecture presented above, one shall find utilization possibilities, and implement them. As an example of the possibilities, a program summing up elements of a vector is presented. The example in Listing 1 has been adapted from the 'asumup' program [13]. The program is written in Y86 assembly language, extended with EMPA metainstructions. The toolchain for EMPA (including the assembler and simulator, as well as some programming methodologies) is described in a separate paper [31]. Since Y86 is a simplified for education version of the widely known Intel x86 processor, the coding can be followed easily. Here only a qualitative description can be given. More programming details are described in [31], the mentioned sample programs, toolchain, their user guides are available from [32].

First the number of arguments and vector address (lines 3-4) are loaded into registers. Then the program clears the sum (line 6), and verifies the number of items (line 8). The kernel of the calculation is in lines 9-15. First the actual vector element is loaded into `%esi`, then it is added (lines 9-10) to the partial sum stored in `%eax`. After this, the address of the actual element is advanced properly (lines 11-12), then

Listing 1. The sum-up routine in traditional coding

```

1 | # This is summing up elements of vector
2 | 0x000: | .pos 0 # Program starts at address 0000
3 | 0x000: 30f206000000 | irmovl $4,%edx # No of items to sum
4 | 0x006: 30f134000000 | irmovl array,%ecx # Array address
5 |
6 | 0x00c: 6300 | xorl %eax,%eax # sum = 0
7 | 0x00e: 6222 | andl %edx,%edx # Set condition codes
8 | 0x010: 7332000000 | je End
9 | 0x015: 506100000000 | Loop: mrmovl (%ecx),%esi # get *Start
10 | 0x01b: 6060 | addl %esi,%eax # add to sum
11 | 0x01d: 30f304000000 | irmovl $4,%ebx #
12 | 0x023: 6031 | addl %ebx,%ecx # Start++
13 | 0x025: 30f3ffffff | irmovl $-1,%ebx #
14 | 0x02b: 6032 | addl %ebx,%edx # Count--
15 | 0x02d: 7415000000 | jne Loop # Stop when 0
16 | 0x032: 00 | End: halt
17 |
18 | # Array of 4 elements
19 | 0x034: | .align 4
20 | 0x034: 0d000000 | array: .long 0xd
21 | 0x038: c0000000 | .long 0xc0
22 | 0x03c: 000b0000 | .long 0xb00
23 | 0x040: 00a00000 | .long 0xa000

```

the loop counter advanced (lines 13-14). Finally, depending on the actual conditions (line 15) the loop returns back to repeat the procedure for the next vector element.

5.1 Eliminating obsolete instructions'(FOR)

As it can be seen from Listing 1, this implementation is rather ineffective: the payload work is done in lines 9-10, while lines 11-15 only serve loop organization. If the SV could take over loop organization, we could reach a performance gain about 3. This needs thinking in terms of EMPA.

Actually, the program fragment needs the address of the current summand and a register to store the partial sum. I.e. we can separate lines 9-10 in a QT and organize the job around this approach. So, lines 9-10 will be executed by the child, on the request from the parent. As mentioned, the parent clones its "glue" to the child, so the address of the array is accessible for the child in `%ecx`, the old partial sum is delivered in `%eax`, and the new partial sum is cloned back to the parent also in `%eax`. The child terminates after executing lines 9-10, allowing the parent to know that one element has been added to the partial sum. So, the job for providing the right address and counting the iterations remains for the parent. Since the parent is only waiting while the child terminates, its arithmetic facilities can be used for this task.

The parent pre-allocates a child for the work. The pre-allocation is needed because the other cores might allocate cores in parallel, so through the preallocation is guaranteed that always will be available core for the iterations. That child core will be allocated by the parent as many times as needed, and so the loop will be jointly executed by the child and the parent. The parent starts executing an iteration (as many times as needed), and waits until the child terminates.

The SV also participates in the game: calculates the address of the vector element for the next iteration and delivers data and signals between parent and child.

The parent writes the current address to its latched 'ForChild' register (see Fig. 2) and it will be latched by the child upon creation into its own the latched 'FromParent' register. The number of the remaining iterations are stored in the latched 'FromChild' in the parent. Before instructing the parent to start a new iteration, SV checks if the content of 'FromChild' is cleared. If not, it instructs the parent to start the next iteration and decrements the latched 'FromChild' in the parent. Since the child can write its latched 'ForParent' register, the content of which is transferred upon termination to the latched 'FromChild' in the parent, the child can **break** the loop.

5.2 Eliminating obsolete stages (SUMUP)

Summing up elements in a vector is a simple and ideal example of processes which cannot be parallelized: in the addition one of the summands is the previous partial sum, so the next summing cannot be started until the previous one is terminated.

The closer look shows that it is because in the frame of executing a single instruction the processors must read the content of a register and write it back, updating it with the new partial sum. We can also see that the partial sum is never used, we are only interested in the final sum. This means, that *if we can find a way in the architecture, which allows a cooperative execution of adding, furthermore a separated readout or the final sum, we can parallelize this strictly sequential process.*

To implement this, needs a bit more functionality from the SV. Suppose we have enough cores which can be allocated, and preallocate them. In this special operating mode which allows opening execution stages for a child, an adder

is prepared in the parent, which on one input receives the latched 'FromChild', and on the other input receives its own output (the previous partial sum). The child receives the current address as described in the previous section, and in this special mode executing `addl` to a special pseudo register means writing to 'ForParent' in the child, which also triggers transferring to 'FromChild' in the parent, i.e. the content read out by the child will be added to the partial sum stored in the parent.

5.3 System services and other uses

Some system services, for example semaphore handling, do not really need all the facilities of the OS, so they can be implemented in alternative way. As our former measurements on soft system [20] proved, such alternative implementation resulted in performance gain about 30, although in that case no context changing was needed. Similar gain can be expected when implementing OS services with EMPA. The gain factor will surely be increased because of the eliminated context change, but the concrete gain will depend on the functionality of the service.

6 PERFORMANCE OF THE ACCELERATOR

As discussed above, using EMPA architecture can distribute the code between PUs, can eliminate obsolete instructions, obsolete execution stages, obsolete context changes, etc. The machine instruction execution remains the same, so to measure performance is not simple at all. Practically, the speedup (the ratio of the execution times) remains the only measurable quantity. Recently, a figure of merit characterizing the *effective parallelization* [33] has been developed, so below that merit will be used to describe the performance gain due to utilizing EMPA, and it will be compared to the traditionally used merit: the speedup divided by the number of cores ($\frac{S}{k}$). The effective parallelization can be derived from the number of PUs k and the measured speedup S as

$$\alpha_{eff} = \frac{k}{k-1} \frac{S-1}{S} \quad (1)$$

The conventional methods of parallelization suffer from inefficiency in using computing power of multiple PUs: because of the presence of the sequential-only part, the more cores are used, the lower is the value of $\frac{S}{k}$, while α_{eff} really describes correctly how effectively the cores are utilized. The 'sumup' program has been tested in three versions: using the conventional programming (i.e. NO EMPA acceleration), replacing "control" instructions with SV activity in the FOR mode, and using SUMUP mode, where (in addition to eliminating control instructions) the cooperation eliminates the unneeded read/write back cycles *within* a machine instruction.

The results (measured using the simulator [32]) are compared in Table 1. for different vector length and different number of cores. The simulator uses arbitrary, but reasonable execution times, expressed in units of the control clock driving the SV. The actual values might change when an electronic version (RC implementation) allows to provide more accurate data.

TABLE 1
Effective parallelization in EMPA processor, in different operating modes

Vector length	Mode of mass proc	Time (clocks)	No of cores (k)	Speedup (s)	$\frac{S}{k}$	α_{eff}
1	NO	52	1	1	1	1
1	FOR	31	2	1.68	0.84	0.81
1	SUMUP	33	2	1.58	0.79	0.73
2	NO	82	1	1	1	1
2	FOR	42	2	1.95	0.98	0.97
2	SUMUP	34	3	2.41	0.80	0.87
4	NO	142	1	1	1	1
4	FOR	64	2	2.22	1.11	1.10
4	SUMUP	36	5	3.94	0.79	0.93
6	NO	202	1	1	1	1
6	FOR	86	2	2.34	1.17	1.15
6	SUMUP	38	7	5.31	0.76	0.95




Fig. 4. Diagram showing data from Table 1: the measurable speedup for two different mass processing methods, in function of the vector length.

6.1 Analyzing speedup results

As Table 1 shows, both the conventional and EMPA execution times increase linearly with the length of the vector. Their intersect and slope values, however, are very different.

The FOR mode of EMPA is nearly 3 times quicker than the conventional method, as some computed control statements are replaced by the much more effective SV functionality. This requires only 2 PUs.

In the SUMUP method, in addition to omitting the computed control machine instructions, even the obsolete fetch, decode, writeback, etc. stages of one instruction execution in the loop kernel are replaced by SV functionality. For the SUMUP mode an extra element increases the execution time by one clock cycle, at the price of utilizing one more PU. This behavior is especially valuable, because using conventional methods of parallelization the algorithm cannot be parallelized at all.

The measured speedup values are derived from a mixture of different types of instructions: both conventional and EMPA code contains both sequential and parallel parts, so despite of the linear increase, the measured speedup will

not linearly depend on the vector length, see Fig. 4. The two speedup values will saturate for high vector lengths at values $\frac{30}{11}$ and 30, respectively.

6.2 Analyzing parallelization efficiency

When eliminating in this very simple loop the control instructions in FOR mode, the $\frac{S}{k}$ values can even be *above* unity. This means *not* a higher PU performance, it is due to the more clever organization of cycles.




Fig. 5. Diagram showing data from Table 1: the core utilization efficiency for two different mass processing methods, in function of the vector length.

In the SUMUP mode, the helper cores are utilized only for a short period of time, so the utilization efficiency is low for short vectors.

Note that since the PUs are put back in the pool, for very long vectors much lower number of cores might be needed. If the compiler can find out the length of processing in that mode (in our example it is 30 clock cycles), it should not allocate more than that number of cores: when the parent needs the 31st core, the 1st core is available again, so the summing can be continued for arbitrary vector length. In calculating the effective utilization of cores using Equ. 1, k should be replaced with $\max_{k \leq k_{eff} \leq 30}$.

The two different points of view of the two merits is best reflected in Fig. 6, where $\frac{S}{k}$ and α_{eff} are shown in function of the vector length, using an EMPA processor in SUMUP mode. Because of the effect of sequential code fragments, both curves start increasing with increasing the number of the cores. As mentioned, in this mass processing mode max. 31 cores (1 parent plus 30 child cores) can be used. There is no sense to use more than 30 child cores: they would need to wait for sending their summand for the parent. Because of this, both the number of threads and the speedup keep raising, while the number of the cores saturates at 31. For short vectors, α_{eff} is relatively low, because the helper cores are used only in a fragment of time. As all the 30 helper cores have "full time job", the α_{eff} dependence saturates at value 1. In contrast, $\frac{S}{k}$ starts to decrease with increasing the number of the cores, and after reaching 30




Fig. 6. Efficiency $\frac{S}{k}$ and α_{eff} for EMPA processor in the SUMUP mode, in function of the vector length.

cores, the speedup continues, but k remains constant, so the dependence turns back and saturates also at value 1, but approaches it much more slowly.

7 CONCLUSIONS

The idea of introducing EMPA in processor technology opens a series of new possibilities. As main accelerator, it allows to turn a many-core processor to an extremely high-power single-core processor. To make an old single-thread program many-core aware, it is enough to recompile the program using an EMPA-aware compiler and run it on an EMPA architecture processor. EMPA uses no hidden PUs: the same cores can be "rented" for normal code execution and out-of-order or speculative evaluation. This means that the superfluous logic [12] concerting and hiding the extra PUs can be omitted, simplifying the internal architecture, reducing the number of transistors and reducing also the power consumption.

For calculational applications, several hundred times higher parallelization can be reached: the compile-time discovery of parallelization possibilities and mixing thread and instruction level parallelism tends to reach the ideal case of "infinite resources" [34]. In addition to that theoretically checked possibility, the SV can more efficiently perform some control functions from loop organization to opening the closed von Neumann execution frames for the helper cores, raising at least one more order of magnitude in the performance.

Using OSs are getting more simple and more effective with using EMPA: no context changing is needed, and the user mode and kernel mode programs can run in (at least partly) parallel. Since QTs are by their nature atomic processing units, a big part of operating systems dealing with semaphores for shared resource usage, scheduling, etc. becomes obsolete (greatly reducing the amount of codes, both writing and testing), and also the built-in synchronization of EMPA can replace those services offered by the OSs.

The real-time characteristics of processors also benefit from EMPA. To service an interrupt, no state saving and restoring is needed, saving memory cycles and code. The program execution will be predictable: the processor need not be stolen from the running main process. The atomic nature of executing QTs will prevent issues like priority inversion, eliminating the need for special protection protocols.

From the point of view of accelerators, an EMPA processor provides a natural interface for linking special accelerators to the processor. Any circuit, being able to handle data and signals shown in Fig. 2 can be linked to an EMPA processor with easy.

REFERENCES

- [1] S. H. Fuller and L. I. Millett, "Computing Performance: Game Over or Next Level?" *Computer*, vol. 44, pp. 31–38, 2011. [Online]. Available: <http://www.cin.ufpe.br/~acm/esd/05688147.pdf>
- [2] IEEE. (2013) IEEE Rebooting Computing. [Online]. Available: <http://rebootingcomputing.ieee.org/>
- [3] D. S. Martin, "Hardware and software techniques for scalable thousand-core systems," Ph.D. dissertation, Stanford University, Berkeley, 2012. [Online]. Available: <http://purl.stanford.edu/mz572jk7876>
- [4] G.-J. V. D. Braak and H. Corporaal, "R-GPU: A Reconfigurable GPU Architecture," *ACM Trans. Archit. Code Optim.*, vol. 13, no. 1, pp. 12:1–12:24, Mar. 2016.
- [5] The Multicore Association, "Software-Hardware Interface for Multi-Many-Core," 2013. [Online]. Available: <http://www.multicore-association.org/workgroup/shim.php>
- [6] J. Yang, H. Cui, J. Wu, Y. Tang, and G. Hu, "Making Parallel Programs Reliable with Stable Multithreading," *Communications of the ACM*, vol. 57, no. 3, pp. 58–69, 2014.
- [7] A. Mahesri, N. J. Wang, and S. J. Patel, "Hardware Support for Software Controlled Multithreading," *SIGARCH Comput. Archit. News*, vol. 35, no. 1, pp. 3–12, Mar. 2007.
- [8] F. Xia, A. Mokhov, A. Yakovlev, A. Iliasov, A. Rafiev, and A. Romanovsky, "Adaptive resource control in multi-core systems," Last accessed Dec. 12, 2015 [Online]. <http://async.org.uk/tech-reports/NCL-EEE-MICRO-TR-2013-183.pdf>, Tech. Rep., 2013.
- [9] Arvind and R. A. Iannucci, "Two fundamental issues in multi-processing," in *4th International DFVLR Seminar on Foundations of Engineering Sciences on Parallel Computing in Science and Engineering*. New York, NY, USA: Springer-Verlag New York, Inc., 1988, pp. 61–88.
- [10] T. David, R. Guerraoui, and V. Trigonakis, "Everything you always wanted to know about synchronization but were afraid to ask," in *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles (SOSP '13)*, 2013, pp. 33–48.
- [11] Saini, S. and Haoqiang Jin and Hood, R. and Barker, D. and more authors, "The impact of hyper-threading on processor resource utilization in production applications," in *High Performance Computing (HiPC), 2011 18th International Conference on*, 2011, pp. 1–10.
- [12] M. Schlansker and B. Rau, "EPIC: Explicitly Parallel Instruction Computing," *Computer*, vol. 33, no. 2, pp. 37–45, Feb. 2000.
- [13] R. E. Bryant and D. R. O'Hallaron, *Computer Systems: A Programmer's Perspective*. Pearson, 2014.
- [14] S. Mahlke, W. Chen, P. Chang, and W.-M. Hwu, "Scalar program performance on multiple-instruction-issue processors with a limited number of registers," in *Proceedings of the Twenty-Fifth Hawaii International Conference on System Sciences*, vol. 1, 1992, pp. 34–44.
- [15] M. R. Williams, "The Origins, Uses, and Fate of the EDVAC," *IEEE Ann. Hist. Comput.*, vol. 15, no. 1, pp. 22–38, Jan. 1993. [Online]. Available: <http://dx.doi.org/10.1109/85.194089>
- [16] A. E. Abdallah, C. Jones, and J. W. Sanders, Eds., *Communicating Sequential Processes. The First 25 Years*. Springer, 2005.
- [17] J. Vég, Z. Bagoly, A. Kicsák, and P. Molnár, "A Multicore-aware Von Neumann Programming Model," in *Proceedings of the 9th International Conference on Software Engineering and Applications (ICSOFT 2014)*, 2014, pp. 150–155. [Online]. Available: <http://www.scitepress.org/DigitalLibrary/Link.aspx?doi=10.5220/2f0005097001500155>
- [18] L. Yavits, A. Morad, and R. Ginosar, "The effect of communication and synchronization on Amdahl's law in multicore systems," *Parallel Computing*, vol. 40, no. 1, pp. 1–16, 2014.
- [19] G. M. Amdahl, "Validity of the Single Processor Approach to Achieving Large-Scale Computing Capabilities," in *AFIPS Conference Proceedings*, vol. 30, 1967, pp. 483–485.
- [20] J. Vég, A. Kicsák, Z. Bagoly, and P. Molnár, "An Alternative Implementation for Accelerating Some Functions of Operating System," in *Proceedings of the 9th International Conference on Software Engineering and Applications (ICSOFT 2014)*, 2014, pp. 494–499. [Online]. Available: <http://www.scitepress.org/DigitalLibrary/PublicationsDetail.aspx?ID=4D%2f9P9PlalB0%3d&t=1>
- [21] H. K.-H. So, "BORPH: An Operating System for FPGA-Based Reconfigurable Computers," Ph.D. dissertation, University of California, Berkeley, 2000.
- [22] Hartenstein, *The Future of Computing, essays in memory of Stamatis Vassiliadis*. Computer Engineering Laboratory, TU Delft, 2007.
- [23] J. Williams, A. D. George, J. Richardson, K. Gosrani, and S. Suresh, "Computational density of fixed and reconfigurable multi-core devices for application acceleration," in *Proceedings of Reconfigurable Systems Summer Institute, Urbana, IL, Jul. 2008.*, 2008.
- [24] Adapteva, "Epiphany Architecture Reference," http://adapteva.com/docs/epiphany_arch_ref.pdf, 2014.
- [25] J. F. Akesson, "Interprocess Communication Utilising Special Purpose Hardware," Ph.D. dissertation, Uppsala University, 2001. [Online]. Available: <http://www.it.uu.se/research/publications/lic/2001-016/2001-016.pdf>
- [26] C. M. Ferreira and A. S. R. Oliveira, *RTOS Hardware Coprocessor Implementation in VHDL*. Intellectual Property / Embedded Systems Conference (IP/ESC), 2009, p. 6. [Online]. Available: http://aveiro.academia.edu/documents/0054/9267/OREK_CoProcessor.pdf
- [27] Cypress, "CY7C026A: 16K x 16 Dual-Port Static RAM," <http://www.cypress.com/documentation/datasheets/cy7c026a-16k-x-16-dual-port-static-ram>, 2015.
- [28] U. Vishkin. (2007) Explicit Multi-Threading (XMT): A PRAM-On-Chip Vision – A Desktop Supercomputer. Last accessed Dec. 12, 2015 [Online]. <http://www.umiacs.umd.edu/users/vishkin/XMT/index.shtml>.
- [29] U. Vishkin, "Using Simple Abstraction to Reinvent Computing for Parallelism," *Communications of the ACM*, vol. 54, no. 1, pp. 75–85, January 2011.
- [30] J. Teich, J. Henkel, A. Herkersdorf, D. Schmitt-Landsiedel, W. Schröder-Preikschat, and G. Snelting, "Invasive Computing: An Overview," in *Multiprocessor System-on-Chip*, M. Hübner and J. Becker, Eds. Springer, 2011, pp. 241–268.
- [31] J. Vég, "Programming facilities for the Explicitly Many-Processor Approach," In *this volume*, p. ??, 2016.
- [32] —, "EMPAthY86: a simulator for EMPA architecture with Y86 cores," 2016. [Online]. Available: <https://github.com/jvegh/EMPAthY86>
- [33] J. Vég, P. Molnár, and J. Vászrhelyi, "A figure of merit for describing the performance of scaling of parallelization," *ArXiv e-prints*, Jun. 2016. [Online]. Available: <http://adsabs.harvard.edu/abs/2016arXiv160602686V>
- [34] A. Nicolau and J. A. Fisher, "Measuring the parallelism available for very long instruction word architectures," *IEEE Transactions on Computers*, vol. C-33, no. 11, pp. 968–976, 1984.



VÉGH, János received his PhD in Physics in 1991, ScD in 2006. He is working in computing since the early 80's, and since 2008 is a full professor in Informatics, currently with University of Miskolc, Hungary. In research, he is looking for ways to prepare more performant and more predictable computing, especially using mainly multi-core processors. He is also dealing with soft processors and hardware-assisted operating systems, reconfigurable and many-processor computing.